arXiv:1609.02175v1 [math.OC] 7 Sep 2016

Tauberian theorems for general iterations of operators: applications to zero-sum stochastic games

Bruno Ziliotto*

Abstract

This paper proves several Tauberian theorems for general iterations of operators, and provides two applications to zero-sum stochastic games where the total payoff is a weighted sum of the stage payoffs. The first application is to provide conditions under which the existence of the asymptotic value implies the convergence of the values of the weighted game, as players get more and more patient. The second application concerns stochastic games with finite state space and action sets. This paper builds a simple class of asymptotically optimal strategies in the weighted game, that at each stage play optimally in a discounted game with a discount factor corresponding to the relative weight of the current stage.

Introduction

Zero-sum stochastic games were introduced by Shapley [17]. In this model, two players repeatedly play a zero-sum game, that depends on an endogenous variable called *state of nature*. At each stage, players choose a pair of actions. Actions and the current state determine both the stage payoff and the distribution according to which the new state is drawn. At the end of the stage, the new state of nature and the pair of actions played are publicly announced to both players.

There are several ways to evaluate the total payoff in a stochastic game. For $n \in \mathbb{N}^*$, the payoff in the $n-$*stage game* is the Cesaro mean $\frac{1}{n}\sum_{m=1}^{n} g_m$, where g_m is the payoff at stage $m \geq 1$. For $\lambda \in (0,1]$, the payoff in the $\lambda-$*discounted game* is the Abel mean $\sum_{m\geq 1} \lambda(1-\lambda)^{m-1} g_m$. More generally, for a sequence of positive weights $(\theta_m)_{m\geq 1}$ summing to 1, the payoff in the θ-weighted game is the weighted sum $\sum_{m\geq 1} \theta_m g_m$. Under mild conditions, the n-stage game, the λ-discounted game and the θ-weighted game have a value, denoted respectively by v_n, v_λ and v_θ.

A huge part of the literature focuses on the existence of the limit of v_n when n goes to infinity, of the limit of v_λ when λ goes to 0, and more generally of the limit of v_θ when $\sup_{m\geq 1} \theta_m$ goes to 0. When the state space and the action sets are finite, Bewley and Kohlberg have proved that (v_n) and (v_λ) converge to the same limit, and a direct consequence of the existence of the uniform value by Mertens and Neyman [10] is that more generally, (v_θ) converges when $(\theta_m)_{m\geq 1}$ is decreasing and θ_1 goes to 0. Without this finiteness assumption, several counterexamples have been pointed out. Vigeral [21] has provided an example of a stochastic game with finite state space and compact action sets in which neither (v_n) nor (v_λ) converges. Ziliotto [24] has provided an example of a stochastic game with compact state space, finite action sets and many other regularity properties, in which (v_n) and (v_λ) fail to converge. Many papers prove convergence of (v_n), (v_λ) and (v_θ) to the same limit in specific

*Paris Dauphine University, Place du Maréchal de Lattre de Tassigny, 75016 Paris, France.
E-mail: ziliotto@math.cnrs.fr

models (see for instance the recent surveys [5], [2] and [18]). A natural question is to ask whether there is a link between the convergence of (v_n), (v_λ) and (v_θ)[*]. Ziliotto [22] has proved that in a very general stochastic game model, with possibly infinitely many states and actions, (v_n) converges uniformly if and only if (v_λ) converges uniformly[†].
This paper aims at generalizing such a result to a more general family of values (v_θ). Note that for recursive games (stochastic games in which the stage payoff is 0 in nonabsorbing states), Li and Venel [7] have proved that if (v_n) or (v_λ) converges uniformly, then (v_θ) converges uniformly when θ is decreasing and θ_1 goes to 0. In a dynamic programming framework (one player), Monderer and Sorin [9] have proved the convergence property, but for a more restricted class of decreasing evaluations.
The contribution of this paper is twofold. In the same stochastic game model as in [22], it provides conditions under which the uniform convergence of (v_n) or (v_λ) implies the uniform convergence of (v_θ). Second, in the case where the state space and the action sets are finite, it proves that the following *discounted strategy* is asymptotically optimal in the θ-weighted game, as θ_1 goes to 0: at each stage $m \geq 1$, play optimally in the discounted game with discount factor $\theta_m/(\sum_{m' \geq m} \theta_{m'})$[‡]. Such a discount factor corresponds to the weight of stage m relative to the weight of future stages. This result is new even when θ is a n-stage evaluation. Finally, this paper provides an example that illustrates the sharpness of the first result.

The proof of the results rely on the operator approach, introduced by Rosenberg and Sorin [16]. This approach builds on the fact that the value of the θ-weighted game satisfies a functional equation, called the Shapley equation (see [17]). The properties of the associated operator can be exploited to infer convergence properties of (v_θ) (see [16]). This paper first proves several Tauberian theorems in an operator setting, and apply them to stochastic games to get the first result. Surprisingly, the proof of the second result also follows from a Tauberian theorem for operators. It is due to the fact that the payoff w_θ guaranteed by a discounted strategy in the θ-weighted game satisfies a functional equation similar to a Shapley equation.

This paper is organized as follows. Sections 1 and 2 state and prove several Tauberian theorems for operators. Section 3 applies these theorems to stochastic games. Section 4 provides an example that shows the sharpness of the results.

1 Operator approach: main results

If $(C, \mathscr{C})$ is a Borel subset of a Polish space, we denote by $\Delta(C)$ the set of probability measures on C, equipped with the weak* topology. In particular, the set $\Delta(\mathbb{N}^*)$ identifies with the set of positive real sequences $(\theta_m)_{m \geq 1}$ such that $\sum_{m \geq 1} \theta_m = 1$. Let $\mathscr{D} := \{\theta \in \Delta(\mathbb{N}^*) \mid \forall\, m \geq 1,\ \theta_{m+1} \leq \theta_m\}$. The sequence $\theta \in \mathbb{R}^{\mathbb{N}^*}$ such that $\theta_m = 0$ for all $m \geq 1$ is denoted by $\mathbf{0}$.

Definition 1.1. Let $\theta \in \Delta(\mathbb{N}^*) \cup \{\mathbf{0}\}$.

- The *shift of* θ is the sequence $\hat{\theta}$ defined by $\hat{\theta} := \mathbf{0}$ if $\theta_1 = 1$, and otherwise

$$\forall\, m \in \mathbb{N}^*, \quad \hat{\theta}_m := \theta_{m+1}(1 - \theta_1)^{-1}.$$

[*] In a continuous-time framework, Oliu-Barton and Vigeral [14] and Khlopin [3, 4] address this question in the optimal control and differential game setting.

[†] Values depend on the initial state, thus (v_n) and (v_λ) map the state space to the reals. Note that if uniform convergence is replaced by pointwise convergence, such a result does not hold, even in the 1-Player case (see Lehrer and Sorin [6]).

[‡] Let us mention that in a recent independent work, Oliu-Barton has obtained similar results, using a different approach. These results will be published in another paper.

- Let $r \in \mathbb{N}^*$. The *r-shift* of θ is the sequence $\theta^r \in \Delta(\mathbb{N}^*)$ defined recursively by $\theta^1 = \theta$ and
$$\theta^r := \widehat{\theta^{r-1}}.$$
Note that if $\theta^r \neq \mathbf{0}$, then for all $m \geq 1$,
$$\theta^r_m = \frac{\theta_{m+r-1}}{\sum_{m' \geq r} \theta_{m'}}.$$

Let $(X, \|.\|)$ be a Banach space, and a map $\Psi : [0,1] \times X \to X$ that satisfies
$$\forall \lambda \in [0,1],\ \forall\ (f,g) \in X^2, \quad \|\Psi(\lambda, f) - \Psi(\lambda, g)\| \leq (1-\lambda)\,\|f - g\|\,, \tag{1.1}$$
and
$$\exists C > 0,\ \forall \lambda \in [0,1],\ \forall\ f \in X, \quad \|\Psi(\lambda, f)\| \leq C\lambda + (1-\lambda)\,\|f\|\,. \tag{1.2}$$

Proposition 1.2. *There exists a unique family $(v_\theta)_{\theta \in \Delta(\mathbb{N}^*) \cup \{\mathbf{0}\}}$ such that: $(v_\theta)_{\theta \in \Delta(\mathbb{N}^*) \cup \{\mathbf{0}\}}$ is bounded, and*
$$\forall\ \theta \in \Delta(\mathbb{N}^*) \cup \{\mathbf{0}\}\,, \quad v_\theta = \Psi(\theta_1, v_{\widehat{\theta}}). \tag{1.3}$$

When θ is taken among decreasing sequences, this proposition is a direct consequence of Neyman and Sorin [13, Subsection 3.2]).

Remark 1.3. A class of operators that play a key role in this paper is the following. Let K be any set, and X be the set of bounded real-valued functions defined on K. Consider two sets S and T, and a family of linear forms $(P_{k,s,t})_{(k,s,t) \in K \times S \times T}$ on X, such that for all (k,s,t), $P_{k,s,t}$ is of norm smaller than 1. Let $g : K \times S \times T \to \mathbb{R}$ be a bounded function. Define $\Psi : [0,1] \times X \to X$ by $\Psi(\lambda, f)(k) := \sup_{s \in S} \inf_{t \in T} \{\lambda g(k,s,t) + (1-\lambda) P_{k,s,t}(f)\}$, for all $f \in X$ and $k \in K$. Note that Ψ satisfies (1.1) and (1.2). This class includes *Shapley operators* (see Neyman and Sorin [12, p.397-415]): this corresponds to the case where K is the state space of some zero-sum stochastic game, S (resp. T) is the set of strategies of Player 1 (resp. 2), k is the current state, and $P_{k,s,t}(f)$ is the expectation of $f(k')$ under strategies s and t, where k' is the state at next stage. Under standard assumptions, for all $\theta \in \Delta(\mathbb{N}^*)$, v_θ is the value of the stochastic game in which the weight of the payoff at stage m is θ_m, for all $m \geq 1$. This fact will be useful in Section 3.

Remark 1.4. Consider a 1-Lipschitz mapping $\Phi : X \to X$ and define $\Psi(\lambda, f) := \lambda \Phi(\lambda^{-1}(1-\lambda) f)$, for all $\lambda \in (0,1]$ and $f \in X$. Extend Ψ to $[0,1] \times X$ by continuity. The operator Ψ satisfies (1.1) and (1.2) for $C = \|\Phi(0)\|$. Thus, the framework considered in this paper is slightly more general than the nonexpansive operator setting that is standard in the literature. This more general setting will be useful in Subsection 3.3.

We emphasize two particular families of evaluations $\theta \in \Delta(\mathbb{N}^*)$:

- For $n \in \mathbb{N}^*$, the element $v_n \in X$ is defined by $v_n := v_{\theta(n)}$, where $\theta(n) \in \mathscr{D}$ is defined in the following way: for all $m \in \mathbb{N}^*$, $\theta(n)_m = 1_{m \leq n} n^{-1}$.
- For $\lambda \in (0,1]$, the element $v_\lambda \in X$ is defined by $v_\lambda := v_{\theta(\lambda)}$, where $\theta(\lambda) \in \mathscr{D}$ is defined in the following way: for all $m \in \mathbb{N}^*$, $\theta(\lambda)_m = \lambda(1-\lambda)^{m-1}$.

By definition, we have
$$v_\lambda = \Psi(\lambda, v_\lambda),$$
and
$$v_n = n^{-1} \Psi\left(\frac{1}{n}, v_{n-1}\right).$$
Some of the results will be stated under the following assumption, which is stronger than (1.2):

Assumption 1. There exists $C > 0$ such that for all $\alpha, \beta \in [0,1]$, for all $\lambda, \lambda' \in [0,1]$, for all $f, g \in X$,

$$\left\|\alpha\Psi(\lambda, f) - \beta\Psi(\lambda', g)\right\| \leq C\left|\alpha\lambda - \beta\lambda'\right| + \left\|\alpha(1-\lambda)f - \beta(1-\lambda')g\right\|.$$

Remark 1.5. Shapley operators (see Remark 1.3) satisfy Assumption 1.

The following proposition stems from the proof of Ziliotto [22, Theorem 1.2] (for further details, see Section 2).

Proposition 1.6. *Under Assumption 1, the following statements are equivalent:*

(a) The sequence $(v_n)_{n\geq 1}$ converges when n goes to infinity.

(b) The mapping $\lambda \to v_\lambda$ has a limit when λ goes to 0.

Moreover, when these statements hold, we have $\lim_{n\to+\infty} v_n = \lim_{\lambda\to 0} v_\lambda$.

When (a), or equivalently (b) is satisfied, the common limit is called *asymptotic value*, and the operator Ψ is said to have an asymptotic value. One of the main goals of this paper is to determine under which conditions on the sequences of weights does the existence of the asymptotic value implies the convergence of (v_θ), as $\sup_{m\geq 1}\theta_m$ goes to 0. This leads to consider the following family of sequences of weights:

Definition 1.7. Let $p \in \mathbb{N}^*$. A sequence $\theta \in \Delta(\mathbb{N}^*)$ is *p-piecewise constant* if there exists $a_1, ..., a_p \in [0,1]$ and $m_1, ..., m_{p+1} \in \mathbb{N}^*$ such that $m_1 = 1$ and for all $p' \in \{1, ..., p\}$, for all $m \in \{m_{p'}, ..., m_{p'+1} - 1\}$, $\theta_m = a_{p'}$, and $\theta_m = 0$ for all $m \geq m_{p+1}$.

For $p \in \mathbb{N}^*$, denote Θ^p the set of sequences that are p-piecewise constant. Note that $\Theta^1 = \{\theta(n), n \in \mathbb{N}^*\}$, and that for $p \geq 2$, Θ^p is not included in $\mathscr{D}$.
For $\theta, \theta' \in \Delta(\mathbb{N}^*)$, define

$$\left\|\theta - \theta'\right\|_1 := \sum_{m\geq 1} |\theta_m - \theta'_m|.$$

For $\Theta \subset \Delta(\mathbb{N}^*)$, denote $D(\theta, \Theta) := \inf_{\theta'\in\Theta} \|\theta - \theta'\|_1$ and
$I_p(\theta) := \max\left\{\sup_{m\geq 1}\theta_m, D(\theta, \Theta^p)\right\}$.
In the stochastic game framework, the quantity $I_p(\theta)$ is a measure of the "impatience" of players. When $I_p(\theta)$ is small, this means first that players are very patient, second that the weight that they put on each stage does not vary too much: it can make at most p significant jumps.

Theorem 1.8. *Suppose that Assumption 1 holds, and that Ψ has an asymptotic value v^*.*
Then for all $\epsilon > 0$, for all $p \in \mathbb{N}^*$, there exists $\alpha > 0$ such that for all $\theta \in \Delta(\mathbb{N}^*)$,

$$I_p(\theta) \leq \alpha \Rightarrow \|v_\theta - v^*\| \leq \epsilon.$$

As an example, let us give a family of evaluations to which the theorem applies.
Let $p \in \mathbb{N}^*$. A sequence $\theta \in \Delta(\mathbb{N}^*)$ is *p-piecewise-discounted* if there exists $a_1, ..., a_p \in \mathbb{R}_+$, $\lambda_1, ..., \lambda_p \in [0,1]$ and $m_1, ..., m_p \in \mathbb{N}^*$ such that $m_1 = 1$ and for all $p' \in \{1, ..., p-1\}$, for all $m \in \{m_{p'}, ..., m_{p'+1} - 1\}$, $\theta_m = a_{p'}(1-\lambda_{p'})^{m-m_{p'}}$, and for all $m \geq m_p$, $\theta_m = a_p(1-\lambda_p)^{m-m_p}$. Denote Θ_d^p the set of sequences that are p-piecewise-discounted. Note that $\Theta^1 = \{\theta(\lambda), \lambda \in (0,1]\}$, and that $\Theta^p \subset \Theta_d^{p+1}$.

Corollary 1.9. *Suppose that Assumption 1 holds, and that Ψ has an asymptotic value v^*. Then for all $\epsilon > 0$, for all $p \in \mathbb{N}^*$, there exists $\alpha > 0$ such that for all $\theta \in \Theta_d^p$,*

$$\sup_{m\geq 1}\theta_m \leq \alpha \Rightarrow \|v_\theta - v^*\| \leq \epsilon.$$

As we shall see in Section 4, this result does not hold when one replaces Θ_d^p by $\mathscr{D}$. For such a result to hold, additional assumptions on the family (v_λ) are required, as shown by the next theorem.

Definition 1.10. The family (v_λ) has *bounded variation* if for all $\epsilon > 0$, there exists $\beta \in (0,1]$ such that for all decreasing sequences $(d_r)_{r\geq 1} \in (0,\beta]^{\mathbb{N}^*}$,

$$\sum_{r=1}^{+\infty} \left\| v_{d_{r+1}} - v_{d_r} \right\| \leq \epsilon.$$

Theorem 1.11. *Assume that one of the following assumptions holds:*

(a) (v_λ) *has bounded variation and Assumption 1 holds.*

(b) There exists $s > 0$ *and* $\beta \in (0,1]$ *such that for all* $\lambda, \lambda' \in (0,\beta]$,

$$\|v_\lambda - v_{\lambda'}\| \leq C \left|\lambda^s - \lambda'^s\right|.$$

Then Ψ *has an asymptotic value* v^**. Moreover, for all* $\epsilon > 0$*, there exists* $\alpha > 0$ *such that for all* $\theta \in \mathscr{D}$*,*

$$\sup_{m\geq 1} \theta_m \leq \alpha \Rightarrow \|v_\theta - v^*\| \leq \epsilon.$$

Remark 1.12. The fact that Assumption (a) or (b) imply the existence of the asymptotic value was already known by [19, Theorem C.8, p.177], or by Proposition 1.6. Note also that when Ψ is the Shapley operator of some stochastic game with finite state space and finite action sets, then (a) and (b) are satisfied (see [1]).

2 Proofs

This section is dedicated to the proofs of the results.

2.1 Iterated operators

Iterating Equation (1.3) in a proper way is one of the key ingredients to the proofs, thus we define the following family of iterated operators:

Definition 2.1. Let $n \in \mathbb{N}$ and $\theta \in \Delta(\mathbb{N}^*) \cup \{\mathbf{0}\}$. The operator $\Psi_\theta^n : X \to X$ is defined recursively by $\Psi_\theta^0 := \mathrm{Id}$, and for $n \geq 1$,

$$\forall\, f \in X, \quad \Psi_\theta^n(f) := \Psi(\theta_1, \Psi_{\hat{\theta}}^{n-1}(f)).$$

When $\theta = \theta(\lambda)$ is a discounted evaluation, we write Ψ_λ^n instead of $\Psi_{\theta(\lambda)}^n$.
The following lemma establishes important properties for the family of operators (Ψ_θ^n).

Lemma 2.2. *Let* $f, g \in X$*,* $\lambda, \lambda' \in [0,1]$*,* $n \in \mathbb{N}^*$*, and* $\theta, \theta' \in \Delta(\mathbb{N}^*) \cup \{\mathbf{0}\}$*. The following assertions hold:*

(i)

$$\|\Psi_\theta^n(f) - \Psi_\theta^n(g)\| \leq \left(\sum_{m\geq n+1} \theta_m \right) \|f - g\|$$

(ii)

$$\|\Psi_\theta^n(f)\| \leq C \sum_{m=1}^{n} \theta_m + \sum_{m \geq n+1} \theta_m \|f\|$$

(iii) Under Assumption 1,

$$\|\Psi_\theta^n(f) - \Psi_{\theta'}^n(f)\| \leq C \sum_{m=1}^{n} \left|\theta_m - \theta'_m\right| + \left| \sum_{m \geq n+1} (\theta_m - \theta'_m) \right| \|f\|.$$

Proof. We introduce the following piece of notation: for $\theta \in \Delta(\mathbb{N}^*) \cup \{\mathbf{0}\}$ and $r \in \mathbb{N}^*$, let

$$\Pi_r(\theta) := \prod_{r'=1}^{r-1} (1 - \theta_1^{r'}) = \sum_{m \geq r} \theta_m.$$

(i) By Equation (1.1), we have

$$\begin{aligned} \|\Psi_\theta^n(f) - \Psi_{\theta'}^n(g)\| &= \left\|\Psi(\theta_1, \Psi_{\hat{\theta}}^{n-1}(f)) - \Psi(\theta_1, \Psi_{\hat{\theta}}^{n-1}(g))\right\| \\ &\leq (1 - \theta_1) \left\|\Psi_{\hat{\theta}}^{n-1}(f) - \Psi_{\hat{\theta}}^{n-1}(g)\right\|. \end{aligned}$$

Iterating this inequality yields

$$\begin{aligned} \|\Psi_\theta^n(f) - \Psi_\theta^n(g)\| &\leq \prod_{m=1}^{n} (1 - \theta_1^m) \|f - g\| \\ &= \left(\sum_{m \geq n+1} \theta_m \right) \|f - g\|. \end{aligned}$$

(ii) Applying Equation (1.2), we get

$$\|\Psi_\theta^n(f)\| \leq C\theta_1 + (1 - \theta_1) \left\|\Psi_{\hat{\theta}}^{n-1}(f)\right\|,$$

and the result follows by induction.

(iii) For $s \in \{0, 1, ..., n\}$, define

$$\Phi(s) := \left\|\Pi_{n-s+1}(\theta)\Psi_{\theta^{n-s+1}}^s(f) - \Pi_{n-s+1}(\theta')\Psi_{\theta'^{n-s+1}}^s(f)\right\|.$$

Let us prove by induction that for all $s \in \{0, 1, ..., n\}$,

$$\Phi(s) \leq C \sum_{m=n-s+1}^{n} \left|\theta_m - \theta'_m\right| + \left| \sum_{m \geq n+1} (\theta_m - \theta'_m) \right| \|f\|. \tag{2.1}$$

We have $\Phi(0) = \left|\Pi_{n+1}(\theta) - \Pi'_{n+1}(\theta')\right| \|f\|$, thus the above inequality holds for $s = 0$. Assume $s \geq 1$. Let $\alpha := \Pi_{n-s+1}(\theta)$, $\beta := \Pi'_{n-s+1}(\theta')$, $\lambda = \theta_1^{n-s+1}$, $\lambda' = \theta_1'^{n-s+1}$, $f := \Psi_{\theta^{n-s+2}}^{s-1}(f)$ and $g := \Psi_{\theta'^{n-s+2}}^{s-1}(f)$. Applying Assumption 1, we get

$$\begin{aligned} \Phi(s) &= \left\|\alpha\Psi(\lambda, f) - \beta\Psi(\lambda', g)\right\| \\ &\leq C \left|\alpha\lambda - \beta\lambda'\right| + \left\|\alpha(1 - \lambda)f - \beta(1 - \lambda')g\right\| \\ &= C \left|\theta_{n-s+1} - \theta'_{n-s+1}\right| + \Phi(s-1). \end{aligned}$$

By induction hypothesis, we deduce that inequality (2.1) holds for all $s \in \{0, 1, ..., n\}$. Taking $s = n$ yields (iii).

□

2.2 Proofs of Propositions 1.2 and 1.6

2.2.1 Proposition 1.2

Let us start with the proof of Proposition 1.2, that is, the existence and uniqueness of the family (v_θ) defined by Equation (1.3). Let $\theta \in \Delta(\mathbb{N}^*) \cup \{\mathbf{0}\}$. For $n \geq 1$, define

$$f_\theta^n := \Psi_\theta^n(0).$$

Note that by Lemma 2.2 (ii), for all $n \in \mathbb{N}^*$, $\|f_\theta^n\| \leq C$. Let $n \geq 1$ and $l \geq 0$. Using Lemma 2.2 (i) and (ii), we get

$$\begin{aligned} \left\| f_\theta^{n+l} - f_\theta^n \right\| &= \left\| \Psi_\theta^n(f_{\theta^{n+1}}^l) - \Psi_\theta^n(0) \right\| \\ &\leq \left(\sum_{m \geq n+1} \theta_m \right) \left\| f_{\theta^{n+1}}^l \right\| \\ &\leq C \sum_{m \geq n+1} \theta_m. \end{aligned}$$

It follows that (v_{θ^n}) is a Cauchy sequence in the Banach space X. Consequently, it converges. Define $v_\theta := \lim_{n \to +\infty} v_{\theta^n}$.

By definition, we have

$$f_\theta^n = \Psi(\theta_1, f_{\hat{\theta}}^{n-1}).$$

Because $\Psi(\theta_1, .)$ is continuous, passing to the limit gives that v_θ satisfies Equation (1.3). Moreover, $(v_\theta)_{\theta \in \Delta(\mathbb{N}^*) \cup \{\mathbf{0}\}}$ is bounded by C.

Now let us prove the uniqueness of the family $(v_\theta)_{\theta \in \Delta(\mathbb{N}^*) \cup \{\mathbf{0}\}}$. Let $(w_\theta)_{\theta \in \Delta(\mathbb{N}^*) \cup \{\mathbf{0}\}}$ satisfying the conditions of Proposition 1.2. Using Lemma 2.2 (i), we deduce that for all $\theta \in \Delta(\mathbb{N}^*)$ and $n \geq 1$,

$$\begin{aligned} \|v_\theta - w_\theta\| &\leq \left(\sum_{m \geq n+1} \theta_m \right) \|v_{\theta^n} - w_{\theta^n}\| \\ &\leq \left(\sum_{m \geq n+1} \theta_m \right) \left(\|v_{\theta^n}\| + \|w_{\theta^n}\| \right). \end{aligned}$$

Because $(v_\theta)_{\theta \in \Delta(\mathbb{N}^*) \cup \{\mathbf{0}\}}$ and $(w_\theta)_{\theta \in \Delta(\mathbb{N}^*) \cup \{\mathbf{0}\}}$ are bounded, taking n to infinity yields that $v_\theta = w_\theta$ for all $\theta \in \Delta(\mathbb{N}^*) \cup \{\mathbf{0}\}$.

2.2.2 Proposition 1.6

As mentioned earlier, this proposition is a direct consequence of the proof in [22, Theorem 1.2]. Indeed, the only difference between Theorem 1.2 in [22] and Proposition 1.6 is that the theorem deals with a more specific class of operators: as in Remark 1.4, the operator Ψ is of the form

$$\forall \lambda \in (0,1] \quad \forall f \in X \quad \Psi(\lambda, f) = \lambda \Phi(\lambda^{-1}(1-\lambda)f),$$

where $\Phi : X \to X$ is a 1-Lipschitz operator. The key point is that the proof of the theorem only relies on the properties of the iterated operator stated in [22, Lemma 1]. Thanks to Lemma 2.2, these properties are also satisfied in our more general framework. Consequently, the same proof applies as well.

2.3 Proof of Theorem 1.8 and Corollary 1.9

In this subsection, Assumption 1 is in force. Let us start with a lemma, that is important to prove both Theorem 1.8 and Corollary 1.9.

Lemma 2.3. *Let $\theta, \theta' \in \Delta(\mathbb{N}^*) \cup \{\mathbf{0}\}$. The following assertion holds:*

$$\|v_\theta - v_{\theta'}\| \leq C \left\|\theta - \theta'\right\|_1 .$$

In particular,

$$\|v_\theta\| \leq C.$$

Proof. Taking $f = 0$ in Lemma 2.2 (iii), we get that for all $n \in \mathbb{N}$,

$$\|\Psi_\theta^n(0) - \Psi_{\theta'}^n(0)\| \leq C \left\|\theta - \theta'\right\|_1 .$$

Taking n to infinity in the above equation yields the first inequality. Taking $\theta' = \mathbf{0}$ in the first inequality implies that $\|v_\theta\| \leq C$ (this is also a consequence of the proof of Proposition 1.2). □

2.3.1 Theorem 1.8

Now let us prove Theorem 1.8. By the previous lemma, it is enough to prove that for any $p \in \mathbb{N}^*$, the following property $\mathscr{A}(p)$ holds:
for all $\epsilon > 0$, there exists $\alpha > 0$ such that for all $\theta \in \Theta^p$,

$$\sup_{m \geq 1} \theta_m \leq \alpha \Rightarrow \|v_\theta - v^*\| \leq \epsilon. \tag{2.2}$$

Let us prove the result by induction on $p \geq 1$. By assumption, (v_n) converges to v^*, thus $\mathscr{A}(1)$ holds. Assume now that $\mathscr{A}(p)$ holds for some $p \geq 1$.
Let $\theta \in \Theta^{p+1}$: there exists $a_1, ..., a_p \in [0,1]$ and $m_1, ..., m_{p+1} \in \mathbb{N}^*$ such that $m_1 = 1$ and for all $p' \in \{1, ..., p\}$, for all $m \in \{m_{p'}, ..., m_{p'+1} - 1\}$, $\theta_m = a_{p'}$, and $\theta_m = 0$ for all $m \geq m_{p+1}$. Thanks to Lemma 2.3, we may assume without loss of generality that $0 < a_1 < 1$ (indeed we can perturbate the weights so that they become positive, and this will not significantly change the value v_θ). Let $\theta' := \theta^{m_2}$. Let n_1 be the integer part of $1/a_1$. We have

$$v_\theta = \Psi_\theta^{m_2 - 1}(v_{\theta'}), \tag{2.3}$$

and

$$v_{n_1} = (n_1)^{-1} \Psi^{m_2 - 1}((n_1 - m_2 + 1) v_{n_1 - m_2 + 1}). \tag{2.4}$$

Moreover,

$$\|v_\theta - v_{n_1}\| \leq \left\|v_\theta - \Psi_\theta^{m_2-1}(v_{n_1-m_2+1})\right\| + \left\|\Psi_\theta^{m_2-1}(v_{n_1-m_2+1}) - v_{n_1}\right\|. \tag{2.5}$$

Equations (2.3), (2.4) and Lemma 2.2 (i) imply that

$$\left\|v_\theta - \Psi_\theta^{m_2-1}(v_{n_1-m_2+1})\right\| \leq (1 - (m_2 - 1)a_1) \|v_{\theta'} - v_{n_1-m_2+1}\| \tag{2.6}$$

and

$$\begin{aligned}
\left\|\Psi_\theta^{m_2-1}(v_{n_1-m_2+1}) - v_{n_1}\right\| &\leq C(m_2 - 1)(1/n_1 - a_1) + (m_2 - 1)(1/n_1 - a_1) \|v_{n_1-m_2+1}\| \\
&\leq 2C(m_2 - 1)(1/n_1 - a_1) \\
&\leq 2C(m_2 - 1)a_1^2/(1 - a_1) \\
&\leq 2Ca_1/(1 - a_1).
\end{aligned} \tag{2.7}$$

Note that the above quantity vanishes as $\sup_{m\geq 1}\theta_m$ goes to 0.
Let $\epsilon > 0$, and let α be as in (2.2). Let us prove that there exists $\alpha' > 0$ such that

$$\sup_{m\geq 1}\theta_m \leq \alpha' \;\Rightarrow\; (1-(m_2-1)a_1)\,\|v_{\theta'} - v_{n_1-m_2+1}\| \leq \epsilon/2.$$

Thanks to Equations (2.5), (2.6) and (2.7), this is enough to prove that $\mathscr{A}(p+1)$ holds. We discriminate between the two following cases:

Case 1 $(m_2-1)a_1 \geq 1-(\epsilon/4C)$

In this case,

$$(1-(m_2-1)a_1)\,\|v_{\theta'} - v_{n_1-m_2+1}\| \quad\leq\quad (\epsilon/4C)\,(\|v_{\theta'}\| + \|v_{n_1-m_2+1}\|) \leq \epsilon/2.$$

Case 2 $(m_2-1)a_1 \leq 1-(\epsilon/4C)$

Because (v_n) converges to v^*, there exists $n_0 \in \mathbb{N}^*$ such that for all $n \geq n_0$,

$$\|v_n - v^*\| \leq \epsilon/4.$$

Note that $n_1 - m_2 + 1 \geq 1/a_1 - 1 - (1-\epsilon/4C)/a_1 = \epsilon/(4Ca_1) - 1$. Moreover,
$\theta'_1 = a_2/(1-(m_2-1)a_1) \leq (4C/\epsilon)\sup_{m\geq 1}\theta_m$. Take $\alpha' = \min(\epsilon/(4C(n_0+1)), \alpha\epsilon/4C)$. Hence

$$\sup_{m\geq 1}\theta_m \leq \alpha' \Rightarrow (1-(m_2-1)a_1)\,\|v_{\theta'} - v_{n_1-m_2+1}\| \leq \epsilon/2,$$

which concludes the proof.

2.3.2 Corollary 1.9

The main idea is to approximate piecewise-discounted evaluations by piecewise-constant evaluations. Indeed, thanks to Lemma 2.3, in order to prove Corollary 1.9, it is enough to prove that for all $p \geq 1$, the following property holds:

$$\forall \epsilon > 0, \quad \exists p' \in \mathbb{N}^*, \quad \forall \theta \in \Theta_d^p, \quad D(\theta, \Theta^{p'}) \leq \epsilon.$$

Let us prove the property for $p = 1$. Let $\epsilon \in (0, 1/10]$, and assume without loss of generality that $1/\epsilon \in \mathbb{N}^*$. Set $p' = \epsilon^{-3}$. Let $\theta \in \Theta_d^1 = \{\theta(\lambda), \lambda \in (0,1]\}$. Let $\lambda \in (0,1]$ such that for all $m \geq 1$, $\theta_m = \lambda(1-\lambda)^{m-1}$. We distinguish between the two following cases:

Case 1 $\lambda \geq \epsilon^3$

Define $\theta \in \Theta^{p'}$ by $\theta'_m = \theta_m$ for $m \in \{1, ..., p'-1\}$, $\theta'_{p'} = (1-\lambda)^{p'-1}$, and $\theta'_m = 0$ for $m \geq p'+1$.
We have $\|\theta - \theta'\|_1 \leq 2(1-\lambda)^{p'-1} \leq (1-\epsilon)^{\epsilon^{-3}-1} \leq \epsilon$.

Case 2 $\lambda < \epsilon^3$

For all $r \geq 1$, define $m_r := r\lfloor \epsilon^2/\lambda \rfloor + 1$. Define $\theta \in \Theta^{p'}$ by $\theta'_m = \theta_{m_r}$ for all $r \in \{1, ..., \epsilon^{-3} - 1\}$ and $m \in \{m_r, ..., m_{r+1} - 1\}$, and $\theta_{m_{\epsilon^{-3}}} := 1 - \sum_{m=1}^{m_{\epsilon^{-3}}} \theta'_m$.
By definition, for all $r \in \{1, ..., \epsilon^{-3} - 1\}$ and $m \in \{m_r, ..., m_{r+1} - 1\}$, we have

$$(1-\lambda)^{\lfloor \epsilon^2/\lambda \rfloor - 1}\theta'_m \leq \theta_m \leq \theta'_m,$$

thus $|\theta_m - \theta'_m| \leq (1 - (1-\lambda)^{\lfloor \epsilon^2/\lambda \rfloor - 1})\theta'_m \leq \epsilon^2\theta'_m$. Moreover, we have

$$\sum_{m \geq m_{\epsilon^{-3}}} \theta_m = (1-\lambda)^{m_{\epsilon^{-3}}} \leq \epsilon^2.$$

Hence $\|\theta - \theta'\|_1 \leq \epsilon^2 + \epsilon^2 \leq \epsilon$ and this proves the property for $p = 1$. By a straightforward induction, the corollary holds.

2.4 Proof of Theorem 1.11

2.4.1 Under Assumption (a)

Let $\epsilon \in (0, 1/10)$ and $\theta \in \mathscr{D}$. Thanks to Lemma 2.3, we assume without loss of generality that for all $m \geq 1$, $\theta_m > 0$. Let $t_1 := 1$, and for $r \geq 1$, define recursively

$$t_{r+1} := \max\left\{m \geq 1 \mid (1-\epsilon)\theta_m^{t_r} \leq \theta_1^{t_r}(1-\theta_1^{t_r})^{m-1} \leq (1+\epsilon)\theta_m^{t_r}\right\} + t_r.$$

with the convention $\max \emptyset = +\infty$ and $t_{r+1} = +\infty$ if $t_r = +\infty$.
In words, between stage 1 and stage $t_r - 1$, the discounted evaluation with discount factor $\theta_1^{t_{r-1}}$ is a good approximation of the evaluation $\theta^{t_{r-1}}$.
Let $\tilde{m} := \max\left\{m \geq 1 \mid \sum_{m' \geq m} \theta_{m'} \geq \epsilon\right\}$ and $\tilde{r} := \max\{r \geq 1 \mid t_r \leq \tilde{m}\}$. Define $m_r := t_r$ for all $r \in \{1, ..., \tilde{r}\}$, and $m_{\tilde{r}+1} := \tilde{m}$. For $r \in \{1, ..., \tilde{r}+1\}$, define $\mu^r := \theta^{m_r}$ and

$$\lambda_r := \mu_1^r = \theta_1^{m_r} = \theta_{m_r}\left(\sum_{m \geq m_r} \theta_m\right)^{-1}.$$

Define $\pi_r \in [0,1]$ by

$$\pi_r := \prod_{r'=1}^{r-1}(1-\lambda_{r'}) = \sum_{m \geq m_r} \theta_m.$$

Lemma 2.4. *The following inequality holds:*

$$\|v_\theta - v_{\lambda_1}\| \leq \sum_{r=1}^{\tilde{r}} \left\|v_{\lambda_r} - v_{\lambda_{r+1}}\right\| + 2C(2\epsilon + \theta_1).$$

Proof. Let $r \in \{1, ..., \tilde{r}\}$. We have

$$v_{\mu^r} = \Psi_{\mu^r}^{m_{r+1}-m_r}(v_{\mu^{r+1}}), \tag{2.8}$$

and

$$v_{\lambda_r} = \Psi_{\lambda_r}^{m_{r+1}-m_r}(v_{\lambda_r}). \tag{2.9}$$

Triangle inequality yields

$$\|v_{\mu^r} - v_{\lambda_r}\| \leq \left\|v_{\mu^r} - \Psi_{\mu^r}^{m_{r+1}-m_r}(v_{\lambda_r})\right\| + \left\|\Psi_{\mu^r}^{m_{r+1}-m_r}(v_{\lambda_r}) - v_{\lambda_r}\right\|. \tag{2.10}$$

Note that

$$\begin{aligned}\left(\sum_{m \geq m_{r+1}-m_r+1} \mu_m^r\right) &= \left(\sum_{m \geq m_{r+1}} \theta_m\right)\left(\sum_{m \geq m_r} \theta_m\right)^{-1} \\ &= \pi_{r+1}/\pi_r.\end{aligned}$$

Thus, by (2.8) and Lemma 2.2 (i), we have

$$\begin{aligned}\left\|v_{\mu^r} - \Psi_{\mu^r}^{m_{r+1}-m_r}(v_{\lambda_r})\right\| &\leq (\pi_{r+1}/\pi_r)\left\|v_{\mu^{r+1}} - v_{\lambda_r}\right\| \\ &\leq (\pi_{r+1}/\pi_r)\left\|v_{\mu^{r+1}} - v_{\lambda_{r+1}}\right\| + (\pi_{r+1}/\pi_r)\left\|v_{\lambda_{r+1}} - v_{\lambda_r}\right\| \end{aligned} \tag{2.11}$$

Equations (2.9) and Lemma 2.2 (iii) yield

$$\begin{aligned}\left\|\Psi_{\mu^r}^{m_{r+1}-m_r}(v_{\lambda_r})-v_{\lambda_r}\right\| &\leq C\sum_{m=1}^{m_{r+1}-m_r}\left|\mu_m^r-\lambda_r(1-\lambda_r)^{m-1}\right| \\ &+\left|\sum_{m\geq m_{r+1}-m_r+1}\mu_m^r-\lambda_r(1-\lambda_r)^{m-1}\right|\|v_{\lambda_r}\| \\ &\leq 2\epsilon C(1-\pi_{r+1}/\pi_r). \qquad (2.12)\end{aligned}$$

The last inequality stems from the definition of (m_r) and the fact that $\|v_{\lambda_r}\|\leq C$. Combining (2.10), (2.11) and (2.12) yield

$$\|v_{\mu^r}-v_{\lambda_r}\|\leq(\pi_{r+1}/\pi_r)\left\|v_{\mu^{r+1}}-v_{\lambda_{r+1}}\right\|+(\pi_{r+1}/\pi_r)\left\|v_{\lambda_{r+1}}-v_{\lambda_r}\right\|+2\epsilon C(1-\pi_{r+1}/\pi_r),$$

thus

$$\pi_r\|v_{\mu^r}-v_{\lambda_r}\|\leq\pi_{r+1}\left\|v_{\mu^{r+1}}-v_{\lambda_{r+1}}\right\|+\pi_{r+1}\left\|v_{\lambda_{r+1}}-v_{\lambda_r}\right\|+2\epsilon C\theta_{m_r}.$$

Summing from $r=1$ to $r=\tilde{r}$ yields

$$\|v_\theta-v_{\lambda_1}\|\leq\pi_{\tilde{r}+1}\|v_{\mu^{r_0}}-v_{\lambda_{\tilde{r}}}\|+\sum_{r=1}^{\tilde{r}}\pi_{r+1}\left\|v_{\lambda_{r+1}}-v_{\lambda_r}\right\|+2\epsilon C.$$

By definition of $\tilde{r}+1$, $\pi_{\tilde{r}+1}\leq\pi_{\tilde{r}+2}+\theta_{\tilde{r}+1}\leq\epsilon+\theta_1$. We deduce that

$$\|v_\theta-v_{\lambda_1}\|\leq\sum_{r=1}^{\tilde{r}}\left\|v_{\lambda_{r+1}}-v_{\lambda_r}\right\|+2C(2\epsilon+\theta_1),$$

and the lemma is proved. □

In Sorin [20, Proposition 4, p.111], by assumption (λ_r) is decreasing, and because (v_λ) has bounded variation, the quantity $\sum_{r=1}^{\tilde{r}}\left\|v_{\lambda_{r+1}}-v_{\lambda_r}\right\|$ can be made as small as desired.
The problem that we face is that in our framework, (λ_r) may not be a decreasing sequence. Nonetheless, it turns out that (λ_r) is "almost" decreasing, in the following sense: the number of integers $r\in\{1,...,\tilde{r}-1\}$ such that $\lambda_r<\lambda_{r+1}$ is at most of order ϵ^{-2}.

Lemma 2.5. *Let $r\in\{1,...,\tilde{r}-1\}$. We have either*

- $\pi_{r+1}\leq(1+\epsilon)^{-1}\pi_r$, *or*
- $\lambda_r\geq\lambda_{r+1}$.

Proof. By definition of m_{r+1}, two cases are possible:

1. $(1-\epsilon)\theta_{m_{r+1}-m_r+1}^{m_r}>\theta_1^{m_r}(1-\theta_1^{m_r})^{m_{r+1}-m_r}$ or
2. $\theta_1^{m_r}(1-\theta_1^{m_r})^{m_{r+1}-m_r}>(1+\epsilon)\theta_{m_{r+1}-m_r+1}^{m_r}$

In the first case, because $(\theta_m)_{m\geq1}$ is decreasing, we have

$$(1-\theta_1^{m_r})^{m_{r+1}-m_r}\leq(1-\epsilon),$$

Moreover, by definition of (m_r), for all $m\in\{1,...,m_{r+1}-m_r\}$,

$$(1+\epsilon)\theta_m^{m_r}\geq\theta_1^{m_r}(1-\theta_1^{m_r})^{m-1}.$$

Combining these two facts, we get

$$\begin{aligned}(1+\epsilon)\sum_{m=1}^{m_{r+1}-m_r}\theta_m^{m_r} &\geq \sum_{m=1}^{m_{r+1}-m_r}\theta_1^{m_r}(1-\theta_1^{m_r})^{m-1}\\ &= 1-(1-\theta_1^{m_r})^{m_{r+1}-m_r}\\ &\geq \epsilon.\end{aligned}$$

Consequently,

$$(1+\epsilon)(1-\pi_{r+1}/\pi_r)\geq\epsilon,$$

thus

$$\pi_{r+1}\leq\frac{1}{1+\epsilon}\pi_r.$$

In the second case, we have

$$\theta_1^{m_r}\geq(1+\epsilon)\theta_{m_{r+1}-m_r+1}^{m_r},$$

and we deduce that

$$\pi_r\lambda_r\geq(1+\epsilon)\pi_{r+1}\lambda_{r+1}.$$

We deduce that either $\pi_r\geq(1+\epsilon)\pi_{r+1}$, or $\lambda_r\geq\lambda_{r+1}$. □

Note that $(1+\epsilon)^{-\epsilon^{-2}-1}<\epsilon$ and $\pi_{\tilde{r}}\geq\epsilon$. Consequently, there exists $k\in\mathbb{N}^*$ satisfying $k\leq\epsilon^{-2}$, and $r_1,...,r_k\in\{1,...,\tilde{r}\}$ such that $r_1=1$, $r_k=\tilde{r}+1$ and for all $k'\in\{1,...,k-1\}$, for all $r\in\{r_{k'},...,r_{k'+1}-1\}$, $\lambda_{r+1}\leq\lambda_r$. To conclude, let $\beta\in(0,1]$ such that for all decreasing sequences $(d_r)_{r\geq1}\in(0,\beta]^{\mathbb{N}^*}$,

$$\sum_{r=1}^{+\infty}\left\|v_{d_{r+1}}-v_{d_r}\right\|\leq\epsilon^3.$$

Set $\alpha:=\beta\epsilon$, and assume that $\theta_1\leq\alpha$. Let $r\in\{1,...,\tilde{r}\}$. Then $\lambda_r\leq\beta\epsilon/\epsilon=\beta$, and thus

$$\sum_{r=1}^{\tilde{r}}\left\|v_{\lambda_{r+1}}-v_{\lambda_r}\right\|\leq k\epsilon^3\leq\epsilon,$$

and the theorem is proved.

2.4.2 Under Assumption (b)

The proof is a simple adaptation of the proof in [23, Theorem 3]. Let $\theta\in\mathscr{D}$. For $r\geq1$, define $\lambda_r:=\theta_1^r=\theta_r/\sum_{m'\geq r}\theta_{m'}$. We have

$$v_{\theta^r}=\Psi_{\theta^r}^1(v_{\theta^{r+1}}),$$

and

$$v_{\lambda_r}=\Psi_{\lambda_r}^1(v_{\lambda_r})=\Psi_{\theta_r}^1(v_{\lambda_r}).$$

These two equations and Lemma 2.2 (i) yield

$$\left\|v_{\theta^r}-v_{\lambda_r}\right\|\leq(1-\lambda_r)\left\|v_{\theta^{r+1}}-v_{\lambda_r}\right\|.$$

This equation is the same as Equation (11) in [23, p.9]. After that, the proof is identical.

3 Applications to zero-sum stochastic games

3.1 Zero-sum stochastic games

Consider a general model of zero-sum stochastic game, as described in Maitra and Parthasarathy [8]. It is described by a state space K, which is a Borel subset of a Polish space, a measurable action set I (resp. J) for Player 1 (resp. 2), that is a Borel subset of a Polish space, a Borel measurable transition $q : K \times I \times J \to \Delta(K)$, and a bounded Borel measurable payoff function $g : K \times I \times J \to \mathbb{R}$.

Given an initial state $k_1 \in K$, the stochastic game Γ^{k_1} that starts in k_1 proceeds as follows. At each stage $m \geq 1$, both players choose simultaneously and independently an action, $i_m \in I$ (resp. $j_m \in J$) for Player 1 (resp. 2). The payoff at stage m is $g_m := g(k_m, i_m, j_m)$. The state k_{m+1} of stage $m+1$ is drawn from the probability distribution $q(k_m, i_m, j_m)$. Then (k_{m+1}, i_m, j_m) is made public to players.

The set of all possible histories before stage m is $H_m := (K \times I \times J)^{m-1} \times K$. A *behavioral strategy* for Player 1 (resp. 2) is a Borel measurable mapping $\sigma : \cup_{m\geq 1} H_m \to \Delta(I)$ (resp. $\tau : \cup_{m\geq 1} H_m \to \Delta(J)$).

A triple $(k_1, \sigma, \tau) \in K \times \Sigma \times \mathscr{T}$ naturally induces a probability measure on $H_\infty := (K \times I \times J)^{\mathbb{N}^*}$, denoted by $\mathbb{P}^{k_1}_{\sigma,\tau}$. Let $\theta \in \Delta(\mathbb{N}^*)$. The $\theta-$*weighted game* $\Gamma^{k_1}_\theta$ is the game defined by its normal form $(\Sigma, \mathscr{T}, \gamma^{k_1}_\theta)$, where

$$\gamma^{k_1}_\theta(\sigma, \tau) := \mathbb{E}^{k_1}_{\sigma,\tau}\left(\sum_{m\geq 1} \theta_m g_m\right).$$

When $\theta = \theta(n)$ for some $n \in \mathbb{N}^*$, the game $\Gamma_n := \Gamma_\theta$ is called the *n-stage game*, and its payoff function is denoted by $\gamma_n := \gamma_\theta$. When $\theta = \theta(\lambda)$ for some $\lambda \in (0, 1]$, the game $\Gamma_\lambda := \Gamma_\theta$ is called the *λ-discounted game*, and its payoff function is denoted by γ_λ. Let $f : K \to \mathbb{R}$ be a bounded Borel measurable function, and $(k, x, y) \in K \times \Delta(I) \times \Delta(J)$. Define

$$\mathbb{E}^k_{x,y}(f) := \int_{(k',i,j)\in K\times I\times J} f(k')dq(k, i, j)(k')dx(i)dy(j)$$

and

$$g(k, x, y) := \int_{(i,j)\in I\times J} g(k, i, j)dx(i)dy(j).$$

We make the following assumption:

Assumption 2. For all $k_1 \in K$ and $\theta \in \Delta(\mathbb{N}^*)$, the game $\Gamma^{k_1}_\theta$ has a value, that is, there exists a real number $v_\theta(k_1)$ such that:

$$v_\theta(k_1) = \sup_{\sigma\in\Sigma} \inf_{\tau\in\mathscr{T}} \gamma^{k_1}_\theta(\sigma, \tau) = \inf_{\tau\in\mathscr{T}} \sup_{\sigma\in\Sigma} \gamma^{k_1}_\theta(\sigma, \tau).$$

Moreover, for all $k_1 \in K$ and $\theta \in \Delta(\mathbb{N}^*)$, the mapping (v_θ) is Borel measurable and satisfies:

$$v_\theta(k_1) = \sup_{x\in\Delta(I)} \inf_{y\in\Delta(J)} \left\{\theta_1 g(k_1, x, y) + (1 - \theta_1)\mathbb{E}^{k_1}_{x,y}(v_{\hat{\theta}})\right\} \tag{3.1}$$

$$= \inf_{y\in\Delta(J)} \sup_{x\in\Delta(I)} \left\{\theta_1 g(k_1, x, y) + (1 - \theta_1)\mathbb{E}^{k_1}_{x,y}(v_{\hat{\theta}})\right\}. \tag{3.2}$$

The game Γ has an *asymptotic value* v^* if both (v_n) and (v_λ) converge uniformly to v^*. Let X be the set of bounded Borel measurable functions from K to $\mathbb{R}$, equipped with the uniform norm, and for all $(f, k) \in X \times K$, define

$$\Psi(\lambda, f)(k) := \sup_{x\in\Delta(I)} \inf_{y\in\Delta(J)} \left\{\lambda g(k, x, y) + (1 - \lambda)\mathbb{E}^k_{x,y}(f)\right\}.$$

We make the following assumption:

Assumption 3. For all $\lambda \in [0,1]$ and $f \in X$, $\Psi(\lambda, f)$ is Borel measurable.

Remark 3.1. When I and J are compact metric spaces, and for all $k \in K$, $q(k,.)$ and $g(k,.)$ are jointly continuous, then Assumptions 2 and 3 hold (see [11, Proposition VII.1.4, p.394]). Maitra and Parthasarathy [8] provide even weaker conditions under which Assumptions 2 and 3 hold in the n stage-case and the λ-discounted case, which generalize immediately to the θ-weighted case.

3.2 Tauberian theorem for value functions

Assume that Assumptions 2 and 3 are in force.
The set X is a Banach space, and Assumption 3 ensures that Ψ is well defined from $[0,1] \times X$ to X. Moreover, Ψ satisfies (1.1) and Assumption 1 for $C := \sup_{(k,i,j) \in K \times I \times J} |g(k,i,j)|$. Thus, Theorems 1.8 and 1.11 apply to Ψ. By Assumption 2, the family of values (v_θ) satisfies equation (1.3). Consequently, the two following theorems hold:

Theorem 3.2. *Assume that Γ has an asymptotic value v^*. Then for all $\epsilon > 0$, for all $p \in \mathbb{N}^*$, there exists $\alpha > 0$ such that for all $\theta \in \Delta(\mathbb{N}^*)$,*

$$I_p(\theta) \leq \alpha \Rightarrow \|v_\theta - v^*\| \leq \epsilon.$$

Theorem 3.3. *Assume that (v_λ) has bounded variation. Then Γ has an asymptotic value v^*. Moreover, for all $\epsilon > 0$, there exists $\alpha > 0$ such that for all $\theta \in \mathscr{D}$,*

$$\sup_{m \geq 1} \theta_m \leq \alpha \Rightarrow \|v_\theta - v^*\| \leq \epsilon.$$

The above theorem was already known. Indeed, if (v_λ) has bounded variation, then the uniform value exists (see [10]). This implies the above theorem (see Theorem 1 and Remark 4 in Neyman and Sorin [13]). Nonetheless, the alternative proof that is provided in this paper is shorter and more elementary.

3.3 Asymptotically optimal strategies in finite stochastic games

Assume that K, I and J are finite. Because (v_λ) can be expanded in Puiseux series (see [1]), there exists $C > 0$ and $s > 0$ such that for all $\lambda, \lambda' \in (0,1]$,

$$\|v_\lambda - v_{\lambda'}\| \leq C \left|\lambda^s - \lambda'^s\right|. \tag{3.3}$$

Let $k \in K$. Denote by $X^*_\lambda(k)$ the set of $x \in \Delta(I)$ such that

$$v_\lambda(k) \quad = \quad \inf_{y \in \Delta(J)} \left\{ \lambda g(k,x,y) + (1-\lambda) \mathbb{E}^k_{x,y}(v_\lambda) \right\}.$$

Denote by $Y^*_\lambda(k)$ the set of $y \in \Delta(J)$ such that

$$v_\lambda(k) \quad = \quad \sup_{x \in \Delta(I)} \left\{ \lambda g(k,x,y) + (1-\lambda) \mathbb{E}^k_{x,y}(v_\lambda) \right\}.$$

Definition 3.4. Let $\theta \in \Delta(\mathbb{N}^*)$. As in the previous section, for $m \in \mathbb{N}^*$, define $\lambda_m := \theta_m / \sum_{m' \geq m} \theta_{m'}$ if $\theta_m \neq 0$, and $\lambda_m := 0$ otherwise. A strategy $\sigma^\theta \in \Sigma$ for Player 1 is *θ-discounted* if for all $m \geq 1$ and $h_m \in H_m$, $\sigma(h_m) \in X^*_{\lambda_m}(k_m)$. Similarly, a strategy $\tau \in \mathscr{T}$ for Player 2 is *θ-discounted* if for all $m \geq 1$ and $h_m \in H_m$, $\tau(h_m) \in Y^*_{\lambda_m}(k_m)$.

The following theorem shows that θ-discounted strategies are asymptotically optimal in Γ_θ, as θ_1 goes to 0.

Theorem 3.5. *For all $\epsilon > 0$, there exists $\alpha > 0$ such that for all $\theta \in \mathscr{D}$, for all $k_1 \in K$, for all θ-discounted strategies $\sigma^\theta \in \Sigma$ and $\tau^\theta \in \mathscr{T}$, for all $\sigma \in \Sigma$ and $\tau \in \mathscr{T}$,*

$$\sup_{m\geq 1} \theta_m = \theta_1 \leq \alpha \Rightarrow \gamma_\theta^{k_1}(\sigma^\theta, \tau) \geq v^*(k_1) - \epsilon \quad \textit{and} \quad \gamma_\theta^{k_1}(\sigma, \tau^\theta) \leq v^*(k_1) + \epsilon.$$

Proof. Let us prove the result for Player 1, the result for Player 2 follows by symmetry. Let X be the set of functions from K to $\mathbb{R}$, equipped with the uniform norm, and for all $(f,k) \in X \times K$, define

$$\Psi(\lambda, f)(k) := \inf_{x \in X_\lambda^*(k)} \inf_{y \in \Delta(J)} \left\{ \lambda g(k,x,y) + (1-\lambda)\mathbb{E}_{x,y}^k(f) \right\}.$$

The operator Ψ satisfies equations (1.1) and (1.2).
For $k \in K$ and $\theta \in \mathscr{D}$, define

$$w_\theta(k) := \inf_{\sigma \ \theta-\text{discounted}} \inf_{\tau \in \mathscr{T}} \gamma_\theta^k(\sigma, \tau).$$

For all $\theta \in \Delta(\mathbb{N}^*)$, we have

$$w_\theta = \Psi(\theta_1, w_\theta).$$

Moreover, for all $\lambda \in (0,1]$, $w_\lambda = v_\lambda$, where v_λ is the value of the λ-discounted game. Applying Theorem 1.11 to Ψ proves the result. □

Remark 3.6. To the best of our knowledge, the above result was not known even when θ is restricted to be a n-stage evaluation.

Remark 3.7. For each $\lambda \in (0,1]$, there exists a stationary optimal strategy $s^*(\lambda)$ (a strategy that depends only on the current state) that is optimal in the game $\Gamma_\lambda(k_1)$, for any $k_1 \in K$ (see [17]). For $\theta \in \mathscr{D}$, consider the following strategy in Γ_θ: at stage m, play the strategy $s^*(\lambda_m)$. By the previous theorem, this strategy is asymptotically optimal, as θ_1 goes to 0. Thus, this theorem provides a simple way to build asymptotically optimal strategies in weighted games.

4 An example

In this section, we provide an example of a 1-Player stochastic game (Markov Decision Process) that has an asymptotic value, and such that there exists $(\theta^N) \in \mathscr{D}^{\mathbb{N}^*}$ satisfying $\lim_{N\to+\infty} \theta_1^N = 0$, but such that (v_{θ^N}) does not converge. For all $N \geq 1$, θ^N is piecewise constant, but the number of pieces grows to infinity as N goes to infinity. This shows first that even if we restrict to decreasing sequences of weights, it is crucial that in Theorem 1.8, the number of pieces is bounded. Second, this example shows that in Theorem 1.11, the bounded variation assumption can not be relaxed. Third, this gives a negative answer to a question raised in Renault [15].

Let Γ be the following MDP: the set space is $K := \{0,1\} \times \mathbb{N}^* \cup \{0^*\}$, and the action set is $I := \{C, Q\}$. The payoff function is equal to 0 in states $\{0\} \times \mathbb{N}^*$, and equal to 1 in states $\{1\} \times \mathbb{N}^*$. The state 0^* is an absorbing state: once in 0^*, the game remains forever in it, and the payoff is 0.

Let $M_0 := 10^{(10^{10})}$. The Dirac measure at $k \in K$ is denoted by δ_k. The transition function q is the following: for all $m \in \mathbb{N}^*$,

$$\begin{aligned}
&q((0,m),C) := (0,m+1)\\
&q((0,m),Q) := [1-(1/\ln(\ln(\ln(m+M_0))))]\cdot\delta_{(1,m^2)} + 1/\ln(\ln(\ln(m+M_0)))\cdot\delta_{0^*}\\
&q((1,m),.) := (1,m-1)\\
&q((1,1),.) := (0,1).
\end{aligned}$$

The idea of the game starting from state $(0,1)$ is the following: the decision-maker plays C for a certain period of time m, and meanwhile gets payoff 0, then he plays Q. At this point, the state is absorbed in 0^* with probability $1/\ln(\ln(\ln(m+m_0)))$, otherwise goes to state $(1,m^2)$. Then, he enjoys a payoff 1 during m^2 stages, and afterwards the game is again in state $(0,1)$. Let us first prove that Γ has an asymptotic value:

Proposition 4.1. *The family (v_n) converges uniformly to the mapping v^* such that $v^*(k)=1$ for all $k \neq 0^*$ and $v^*(0^*)=0$. Consequently, Γ has an asymptotic value.*

Proof. For $n \in \mathbb{N}^*$, consider the following stationary strategy in the n-stage game: play C until the state is $(0,m)$ with $m \geq n^{1/2}$, then play Q, and afterwards play C forever. Under such a strategy, whatever be the initial state, the expected payoff is larger than

$$(1-n^{-1/2}-n^{-1})\left[1-1/\ln(\ln(\ln(n^{1/2}+M_0)))\right].$$

This quantity goes to 1 as n goes to infinity. □

Proposition 4.2. *There exists a sequence $(\theta^N) \in \mathscr{D}^{\mathbb{N}^*}$ such that (θ^N_1) goes to 0 but the real sequence $(v_{\theta^N}(0,1))$ converges to 0.*

Remark 4.3. This result answers negatively the open question mentioned in Renault [15, Section 5]: indeed, the total variation of θ^N vanishes as N goes to infinity, but (v_{θ^N}) does not converge to the asymptotic value.

Proof. For $N \in \mathbb{N}^*$ and $r \in \{1,...,N+1\}$, define $m^N_r := \sum_{r'=1}^{r} N^{3^{r'}} - N^3 + 1$. Consider $\theta^N \in \Delta(\mathbb{N}^*)$ defined by

$$\theta^N_m := \begin{cases} N^{-3^{r+1}-1} & \text{if } r \in \{1,...,N\} \text{ and } m^N_r \leq m < m^N_{r+1},\\ 0 & \text{if } m \geq m^N_{N+1}. \end{cases}$$

Consider $\Gamma_{\theta^N}(0,1)$, and fix a strategy for the decision-maker. Let $r \in \{1,2,...,N\}$. We distinguish between the two following cases:

Case 1 The decision-maker does not play Q in a state $(0,m) \in \{0\}\times\mathbb{N}^*$ between stages m_r and $m_{r+1}-1$

In this case, the only way that his total payoff between stages m_r and $m_{r+1}-1$ is positive is that $k_{m_r} = (1,\tilde{m})$, for some $\tilde{m} \in \mathbb{N}^*$. By definition of the transition, we have

$$\tilde{m} \leq \left(\sum_{r'=1}^{r} N^{3^{r'}}\right)^2 \leq N^{3^{r+1}-1}.$$

We deduce that the total payoff between stages m_r and $m_{r+1}-1$ is smaller than $\tilde{m}N^{-3^{r+1}-1} \leq (1/N)^2$.

Case 2 The decision-maker plays Q in a state $(0,m) \in \{0\} \times \mathbb{N}^*$ between stages m_r and $m_{r+1}-1$

By definition of the transition, we have $m \leq \sum_{r'=1}^{r} N^{3^{r'}} \leq N^{3^{r+1}} \leq N^{3^{N+1}}$, thus $1/\ln(\ln(\ln(m+M_0))) \geq (10\ln(N))^{-1} \geq N^{-1/4}$.

Let us now finish the proof. Assume $N \geq 10$. For $m \in \mathbb{N}^*$, let $p(m)$ be the probability that the state is 0^* at stage m. Assume that under the strategy of the decision-maker, $p(m_{N+1}^N) \geq N^{-1}$. Because $(1-N^{-1/4})^{\sqrt{N}} < N^{-1}$, this implies that the number of r that satisfy Case 2 is smaller than $\sqrt{N}$. Thus, the total payoff in $\Gamma_\theta(0,1)$ is smaller than $N^{-1}+N^{-1/2}$.

Assume now that under the strategy of the decision-maker, $p(m_{N+1}^N) < N^{-1}$. Let m be the first stage such that $p(m) \leq N^{-1}$. Because $(1-N^{-1/4})^{\sqrt{N}} < N^{-1}$, the number of r that satisfy Case 2 and such that $m_r \leq m$ is smaller than $\sqrt{N}$. Like before, we deduce that the total payoff between stages 1 and m is smaller than $N^{-1}+N^{-1/2}$, and afterwards it is smaller than N^{-1} by definition of m. This proves that $\lim_{N\to+\infty} v_{\theta^N}(0,1) = 0$. □

Acknowledgments

The author is very grateful to Miquel Oliu-Barton, Jérôme Renault, Sylvain Sorin and Guillaume Vigeral for helpful discussions.

References

[1] T. Bewley and E. Kohlberg. The asymptotic theory of stochastic games. *Mathematics of Operations Research*, 1(3):197–208, 1976.

[2] A. Jaskiewicz and A.S. Nowak. Zero-sum stochastic games. *preprint*, 2016.

[3] D. Khlopin. On uniform tauberian theorems for dynamic games. *preprint arXiv:1412.7331*, 2014.

[4] D. Khlopin. On asymptotic value for dynamic games with saddle point. *preprint arXiv:1501.06933*, 2015.

[5] Rida Laraki and Sylvain Sorin. Advances in zero-sum dynamic games. *Handbook of Game Theory*, 4:27–95, 2015.

[6] E. Lehrer and S. Sorin. A uniform tauberian theorem in dynamic programming. *Mathematics of Operations Research*, 17(2):303–307, 1992.

[7] X. Li and X. Venel. Recursive games: uniform value, tauberian theorem and the mertens conjecture "maxmin=lim v(n)= lim v(lambda)". *Published online in International Journal of Game Theory*, 2015.

[8] A Maitra and T Parthasarathy. On stochastic games. *Journal of Optimization Theory and Applications*, 5(4):289–300, 1970.

[9] A Maitra and W Sudderth. Borel stochastic games with lim sup payoff. *The Annals of Probability*, pages 861–885, 1993.

[10] J.F. Mertens and A. Neyman. Stochastic games. *International Journal of Game Theory*, 10(2):53–66, 1981.

[11] J.F. Mertens, S. Sorin, and S. Zamir. *Repeated games*. CORE DP 9420-22, 1994.

[12] A. Neyman and S. Sorin. *Stochastic games and applications*. Kluwer Academic Publishers, 2003.

[13] A. Neyman and S. Sorin. Repeated games with public uncertain duration process. *International Journal of Game Theory*, 39(1):29–52, 2010.

[14] M. Oliu-Barton and G. Vigeral. A uniform tauberian theorem in optimal control. In *Advances in Dynamic Games*, pages 199–215. Birkhaüser, 2013.

[15] J. Renault. General limit value in dynamic programming. *Journal of Dynamics and Games*, 1(3):471–484, 2014.

[16] D. Rosenberg and S. Sorin. An operator approach to zero-sum repeated games. *Israel Journal of Mathematics*, 121(1):221–246, 2001.

[17] L.S. Shapley. Stochastic games. *Proceedings of the National Academy of Sciences of the United States of America*, 39(10):1095–1100, 1953.

[18] E. Solan and B. Ziliotto. Stochastic games with signals. In *Advances in Dynamic and Evolutionary Games*, pages 77–94. Springer, 2016.

[19] S. Sorin. *A first course on zero-sum repeated games*, volume 37. Mathématiques et Applications, Springer, 2002.

[20] S. Sorin. Asymptotic properties of monotonic nonexpansive mappings. *Discrete Event Dynamic Systems*, 14(1):109–122, 2004.

[21] G. Vigeral. A zero-sum stochastic game with compact action sets and no asymptotic value. *Dynamic Games and Applications*, 3(2):172–186, 2013.

[22] B. Ziliotto. A tauberian theorem for nonexpansive operators and applications to zero-sum stochastic games. *preprint arXiv:1501.06525, to appear in Mathematics of Operations Research*, 2015.

[23] B. Ziliotto. General limit value in zero-sum stochastic games. *International Journal of Game Theory*, 45(1):353–374, 2016.

[24] B. Ziliotto. Zero-sum repeated games: counterexamples to the existence of the asymptotic value and the conjecture maxmin= lim v (n). *The Annals of Probability*, 44(2):1107–1133, 2016.